European Centre
for Medium Range Weather Forecasts

ECMWF NEWSLETTER

Shinfield Park, Reading, Berkshire RG2 9AX, England

Reading (0734) 85411 Telex 847908

Number 4 - August 1980

Number 4 - August 1980

IN THIS ISSUE:

* NOTE : These articles directly concern the computer service, we recommend computer users read them all.

COVER: General view of the Computer Hall at Shinfield Park

This Newsletter is edited and produced by User Support.

The next issue will appear in October.

THE SECOND SESSION OF THE TECHNICAL ADVISORY COMMITTEE

The ECMWF Technical Advisory Committee held its second session at the ECMWF Headquarters, Shinfield 3 - 6 June 1980, immediately after the first annual forecaster's meeting, which is reported elsewhere (p. 2 of this Newsletter). Mr. J. Lepas (France) was elected Chairman and Mr. W. Wann (Ireland) Vice-Chairman of the Committee.

The most important item on the agenda of the session was the consideration of the Centre's proposals regarding its technical activity and development of technical facilities contained in the Draft Budget for 1981 and the four-year Plan of Activities for 1981 - 1984, including the overall strategy for provision of computer resources for the period 1981 - 1984 and beyond. The Committee basically agreed with the Centre's plans but it is not possible to give details at the present time, as the Draft Budget for 1981 and Plan of Activities have also to be considered at the next two sessions of the Finance Committee and finally by the Council at its twelfth session on 20 - 21 November 1980.

On the basis of the report by the Chairman of the forecasters' meeting, Dr. R. Berggren of Sweden, and a presentation by Mr. A. Lange (Finland) on some preliminary tentative results of the NWP data study/intercomparison project initiated by the WMO CAS Working Group on Weather Prediction Research and being carried out by the Finnish Meteorological Institute, the Technical Advisory Committee agreed that the Centre's model outputs are on average the most accurate available. More specifically, the views expressed by the forecaster's meeting on the quality of the Centre's products (see page 2) were endorsed.

The Committee also considered the possibility of changing the Centre's daily operational schedule. At present, the forecast is run overnight based on the virtually full set of 12Z data, and the results to 7 days are available to Member States over the telecommunications network by around 02 or 03Z the following morning. An alternative schedule would be to run the forecast from about 03Z to 06Z from an initial state taking account of 18Z and early 00Z data. In this case, the results to 7 days would be available to Member States by 06 to 07Z. However, by including the later data, the range of usefulness of the forecasts might be expected to be extended by several hours. In the course of its discussion the Committee reviewed the results of experiments comparing the operational forecasts as currently made with test forecasts starting from a 00Z initial state based on the early 00Z data. The results of these tests were not firmly conclusive and the Committee requested that further experiments should be undertaken in order to determine the impact on the quality of the medium range forecasts. In deciding upon the daily operational routine at the Centre, the differing requirements of users in Member States have also to be taken into account (i.e. the dead-line by which Member States should receive the results and process them if necessary to be available in time for forecasters to make use of the products). The Committee agreed to consider this matter again at its next session. In view of the Committee deliberations, no major change will be made in the Centre's daily operational routine before, say, 1 April 1981, and only then if the continued experiments show that the forecasts can be significantly improved by a modified schedule.

It is not possible fully to report here all the results of the four day Committee session. More comprehensive and formal summaries have been prepared and distributed to staff directly concerned in the Centre and Member States (e.g. "Actions of the Technical Advisory Committee at its second session" ECMWF/TAC/D(80)1.) The Committee's recommendations will go forward to Council for final approval.

Overall, the second session of the Technical Advisory Committee was highly constructive and successful, and as might be expected from an "Advisory Committee", much useful advice, and many comments and recommendations regarding aspects of the Centre's operational activity resulted. The third session of the Technical Advisory Committee will probably take place in February 1981.

- Daniel Söderman

* * * * * * * * * *

MEETING OF FORECASTERS AT SHINFIELD PARK, 2-3 JUNE 1980

At its first session held at Shinfield Park 4 - 6 September 1979, the Technical Advisory Committee urged that an annual meeting be held between forecasters who directly use the Centre's products to interchange views on synoptic quality of the products. Accordingly, a meeting to discuss the synoptic quality of ECMWF products was held at Shinfield Park on June 2 and on the morning of June 3 1980 with the primary object of the meeting being to gather the subjective feelings from the users in the Member States who are studying or interpreting the Centre's forecasts and to hear their synoptic experience and comments regarding the behaviour of the Centre's model. The meeting was attended by forecasters or representatives from ten Member States.

The first day of the meeting was devoted to presentation of papers assessing the ECMWF forecasts. Evaluation of the forecasts were in the main subjective, giving generally the feeling of the forecasters in the Member States of the degree of usefulness of the forecasts as far as seven days. There were some reports also of quantitative or objective assessment of the Centre's products, and also comparison of Centre forecasts with other forecasts currently available in the Member States. Systematic errors in the forecasts (e.g. a reduction in the geopotential height of the mid-tropospheric pressure levels during the forecast, errors in forecasting Mediterranean depressions, 850 mb temperatures being forecast sometimes too low over Southern Europe, occasional inconsistencies between surface and upper-level features) were noted by some of the speakers.

The second day of the meeting was devoted to general discussion and the formulation of an overall view of the quality and characteristics of the Centre's forecasts, based on the various presentations of the previous day. The conclusions of the meeting were:

i) subjective evaluation of the ECMWF forecasts using conventional synoptic verification techniques for specific areas of interest reached on average the mid-point of the quality scale at between 4 and 5 days of the forecast, the scores for 500 mb tending to be somewhat better than those for 1000 mb;

ii) as an average of Member States evaluations there is a high percentage of usable or better forecasts of the 500 mb geopotential field at 5 and 6 days (60% at 5 days, 40% at 6 days);

iii) the ECMWF forecasts are also the best on average of the available 1 to 3 day forecasts.

The meeting also agreed on the need for, and value of, the Member States continuing their subjective evaluation of the Centre products, noting that different Member States had different emphasis in evaluation of the products, e.g. different synoptic features are important for different forecast centres. The results of these evaluations should be considered at a future forecasters' meeting. There was some discussion as to who were the main users of the medium-range forecast products within the Member States, in such areas as agriculture, marine, construction and energy industries. The meeting also believed it would be valuable to establish a meteorological contact point in the Member States to facilitate liaison, and exchange of ideas of meteorological and operational interest in relation to the Centre's operational products.

It is planned to hold another Forecasters Meeting in 1981. It is hoped that a longer period of time can be allotted for the meeting so that as well as interchange of views on the synoptic quality of the products, including model characteristics, practical exploitation of the model results including interpretation techniques, the value of the forecasts to customers and ways of presenting the forecasts, the standardisation of subjective verification techniques and use of suitable objective verifications could also be discussed.

- Austin Woods

* *

POST-PROCESSING

This is an article in the series describing the various subsystems that make up the complete operational suite.

Post-processing handles the organisation and distribution of analysis and forecast results from an ECMWF operational forecast run. To hold the results in suitable forms, a number of databases are set up on the Cyber:

i) a global fields database (GFDB) to feed the archiving;
ii) a polar stereographic database (PSDB) which, with the GFDB, feeds the operational watch;
iii) a dissemination database (DDB) holding the products for dissemination via the communications computer to Member States.

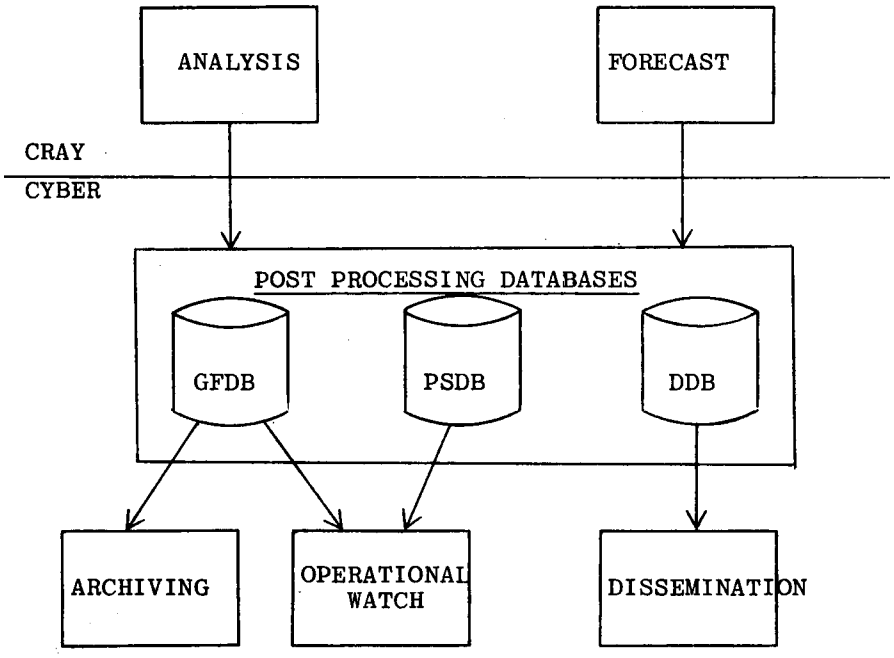

At regular intervals during an operational run, jobs are started in the Cray to create history files of field data covering the globe at standard pressure levels ranging from the surface up to 30mb. The history files are disposed to the Cyber where they are transformed and stored in databases by a reception and transformation program. Each timestep disposes three types of data in separate history files:

i) physics fields of mainly surface parameters on the model latitude - longitude grid (1.875^{o} mesh);
ii) upper air fields on a regular latitude-longitude grid (1.5^{o} mesh);
iii) the same upper air fields as spectral co-efficients.

The data types are handled in independent substeps which can proceed in parallel and asynchronously. This logical division into substeps is maintained in the databases and also in the backup and regeneration procedures since it means that run failures or computer file problems can often be localised and corrected by reprocessing a single substep.

Timesteps at which history files are generated

ANALYSIS

day	hour	data
D-1	18Z	initialised analysis, 6 hour forecast to next timestep
D	00Z	initialised analysis, 6 hour forecast to next timestep
D	06Z	initialised analysis, 6 hour forecast to next timestep
D	12Z	uninitialised analysis, initialised analysis

FORECAST

day	hour	data
D	18Z	6 hour forecast
(every 6 hours)		
D+2	00Z	36 hour forecast
(every 12 hours)		
D+10	12Z	240 hour forecast

Day D is the date of the operational run

Large amounts of information are passed in the history files. Each of the 31 timesteps generates about 3 million numbers which must be transferred over the Cray-Cyber link and stored on Cyber mass-storage devices. To reduce the transfer times and the storage volume, the numbers are packed 4 values to a computer word in the Cray and kept in this condensed form during processing and in the databases. History files are kept on Cyber discs as a backup of the global fields database which itself serves as a backup for the other two databases. To maintain this backup capability, history files and databases are spread over five separate Cyber disc units; and the whole collection of fields is kept online until purged by the next operational run. The purging process is split into two to reduce the number of discs on-line in the early part of the run; the first post-processing family purges files up to D+3, and the first job after D+3 purges the remainder up to D+10. History files are purged by a standalone program running at the beginning of the Suite.

A set of history files for each timestep is processed as a family of jobs started and controlled by the Supervisor-Monitor-Scheduler program (SMS) handling the operational run. After running an initialisation job in a timestep, SMS starts three substeps to run in parallel. Each substep consists of the same reception and transformation program operating on a different data type using time and data type parameters supplied by SMS. A requirements file built up from the dissemination products catalog and the expected automatic display requirements controls the build-up of the PSDB and DDB, and a control file and its backup are updated as each database file is completed in a substep. Substeps can be restarted individually, in which case processing restarts after the most recently completed pressure level. Most fields pass through memory only once with all the derived fields being created at one time; this reduces input-output overheads but makes it necessary to operate on fields in packed format throughout. Every access to the databases is made through general fields database access software (FIDABA). This access software consists of FORTRAN-callable subroutines and uses standard Cyber Record Manager and NOS/BE permanent file functions. Similar file structures are used in all three databases resulting in many routines being shared, and allowing GFDB, PSDB and DDB files from the same substep to be created in parallel.

When fields have been inserted in the databases, they are available for plotting and dissemination. At main analysis and forecast verifying times such as midnight and midday, SMS also launches automatic plotting jobs within a timestep family triggered by completion of the substeps. Similarly at midday timesteps throughout the operational suite a job is run to interrogate the DDB and send files of products to the telecommunications computer. FIDABA serves as the database access mechanism in all cases.

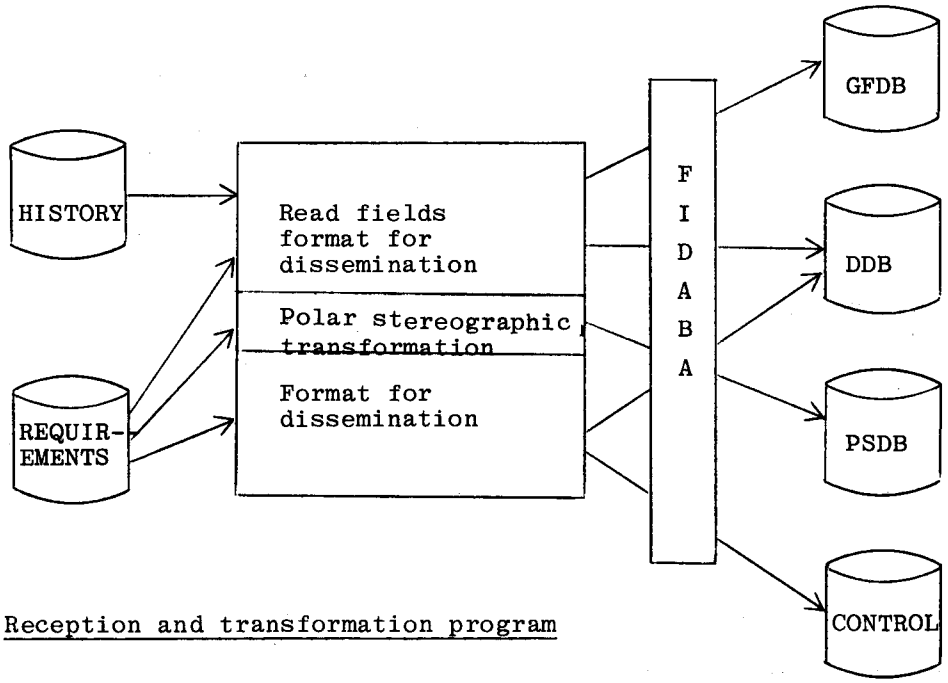

Reception and transformation program

The global fields database comprises all the history file data indexed so that individual fields can be retrieved at random via FIDABA. The surface fields and the upper air spectral coefficients are archived from the GFDB before the next operational run when all the databases are purged.

The polar stereographic database consists of a selection of about 1400 fields transformed to a regular rectangular grid on a northern or southern hemisphere polar stereographic projection, with the grid interval set to 150km true at latitude 60°. The fields are individually accessible using FIDABA and provide field data for the automatic display and visualisation programs in the Operational Watch. The PSDB fields will eventually also serve as backup for dissemination products based on polar stereographic grids.

In the dissemination database, there are currently about 3000 dissemination products built up from latitude-longitude or polar stereographic grids over standard areas selected to suit Member States'needs. The products are coded into three possible formats:

i) ECMWF bit code - a condensed computer-to-computer oriented code for medium speed lines;
ii) WMO GRID code, CCITT alphabet 2 - for low speed lines;
iii) WMO GRID code, CCITT alphabet 5 - for low or medium speed lines.

Except for insertion of bulletin sequence number, DDB products are ready for collection into files for individual Member States and immediate despatch to the telecommunications computer. The products remain on-line available for retransmission until purged at the start of the next operational run.

- John Chambers

* * * * * * * * * * * * *

ECMWF METEOROLOGICAL PUBLICATIONS JUNE - JULY 1980

Technical Report No. 18 : Confidence limits for verification and energetics studies.

Workshop on stochastic, dynamic forecasting, held 17-19 October 1979.

Technical Memorandum No. 17 : On some predictability problems.

Technical Memorandum No. 18 : Some remarks on the effects of the Qinghai-Tibet plateau on weather developments.

* * * * * * * * * * * * *

VACANCIES AT ECMWF

There are the following vacancies at ECMWF Headquarters at Shinfield, near Reading. Remuneration is commensurate with those of International Organisations. For both posts, fluency in at least one of the working languages of the Centre is required; a working knowledge of the other working languages would be an advantage. (The working languages are French, English and German). For further information contact Personnel Section.

POST: SENIOR SCIENTIST (R202)

FUNCTION: To carry out research into methods for the assimilation of meteorological observations into high resolution numerical models and into methods for the initialisation of numerical models.

QUALIFICATIONS: A university education (Ph.D) or equivalent and several years post graduate experience in dynamical meteorology and related areas of meteorology. Practical experience in numerical weather prediction in the fields of data assimilation, initialisation and objective analysis would be an advantage.

STARTING DATE: 1 November 1980 or as soon as possible thereafter.

POST: JUNIOR SCIENTIST/PROGRAMMER (R604)

FUNCTION: To carry out scientific programming and system analysis for the Centre's numerical models. This will mean, in particular, to design and maintain the diagnostic packages for the ECMWF global forecasting models, the maintenance of the research forecasting system and program libraries.

Additional functions will include liaison with the Operations Department to implement changes to the operational forecasting system.

QUALIFICATIONS: A university education or equivalent in a scientific or mathematical discipline and several years programming experience, preferably as a member of a team.

Background in meteorology or another allied science is desirable as well as practical experience in large-scale scientific computing.

STARTING DATE: 1 September 1980 or as soon as possible thereafter.

* * * * * * * * * * * * * * *

CHANGES TO THE CENTRE'S OPERATIONAL FORECASTING ROUTINE AND SCHEDULE

The routine operational medium-range forecasting at the Centre began on 1 August 1979 and since that time forecasts have been produced five days a week. In accord with the overall long-term plans for the development of the Centre's operational activity, forecasts will be carried out for the full seven days a week with effect from 1 August 1980.

At the same time the cut-off time for data used in determining the 12Z initial state for the forecast will be changed from 1745Z to 2045Z. This change is to allow a small, but important, amount of extra data to be included in the 12Z data assimilation (mainly from the Southern Hemisphere and SHIP reports). With data now being received in batches from the United Kingdom Meteorological Office via a medium-speed connection, the later cut-off is practical and there is not a similar delay to the Centre's daily operational routine. Instead of beginning at around 2100Z as at present, the operational forecast will begin at about 2230Z, and thus will normally end about 1½ hours later than at present, i.e. around 0230Z. Correspondingly, the dissemination of results from the Centre to Member States over the telecommunications network will be about 1½ hours later than now.

- Roger Newson

* * * * * * * * * * * * * * * * *

NEW HEAD OF COMPUTER DIVISION

Following Rob Brinkhuysen's resignation, I joined ECMWF on May 19th, 1980, first as a consultant and then on July 7th, 1980 as the new Head of Computer Division.

After studying mathematics, physics, statistics and philosophy at the Universität Tübingen, I was granted the degree of 'Diplom-Mathematiker' in 1969. My thesis was on mathematical theorem proving by computer. The following years I spent first doing some research work on artificial intelligence, then working in the Computing Centre of the Universität Tübingen, becoming in 1972 the Head of the Systems Group there. The Centre runs a TR440 and a CDC 3300 computer system. At the same time, I was lecturer for computer science at the University. In 1978, I was granted a two year fellowship by CERN. There, I worked on the IBM systems for providing small computer support via CERNET.

From 1975 - 1977 I was chairman of the 3000L Main Frame Group and thus a member of the Board of Directors of ECODU.

- Geerd R. Hoffmann

* * * * * * * * * *

* ACCOUNT PROJECT VALIDATION

Recently, the initial stage of project validation was introduced. The messages produced by validation indicate the validity of the project and of the priority specified, plus a detailed breakdown of the allocations and usage if the project is valid.

Allocations are given for Cyber, Cray, and Cyber permanent file resources. The Cyber and Cray allocations are broken down into allocations for priority groups. The groups are BASE (priority 0), LOW (priority 2), NORMAL (priority 4), HIGH (priority 6), MET.OP (priorities 10 - 13) and OPERATIONS (priorities 14 - 17), as described in ECMWF Newsletter No. 2.

The default allocations to each group are 45% to LOW, 45% to NORMAL, 10% to HIGH and none to any other group (except BASE which is unlimited).

Usage of Cyber and Cray resources is also collected by the priority group in which each job runs. However, owing to a system design, if the operators reset the priority of a job in the input queue (an exceptional occurrence) that job will be charged at the priority to which it is reset. Also, any jobs at priorities 12, 13 will be charged at OPERATIONS priority. This incompatibility will be removed in the near future. This does not affect the priority used for validating the project.

Permanent file usage is also recorded for each project. It may be noted that 1 RB (Request Block) is (currently) 56 PRUs, or 3.584 Kilo-words.

- Tony Stormer

* * * * * * * * * *

REQUESTING A FILE ON CYBER MASS STORAGE

Normally, it is not necessary to request a file to be resident on mass storage, unless a particular attribute is needed, e.g.

```
REQUEST,FILE,*Q.        REQUEST FILE ON QUEUE DEVICE
REQUEST,FILE,*PF.       REQUEST FILE ON PF DEVICE
```

However, there may be certain applications borrowed from other installations which make requests for files on a certain type of mass storage, e.g.

```
REQUEST,FILE,*AZ.       REQUEST FILE ON 844-41
```

Such requests should be modified to be device independent, as this autumn we shall be running with both 844-41 (AZ) and 885 (AJ) RMS. Therefore, all general file requests should be of the form:

```
REQUEST,FILE,*A*.       REQUEST FILE ON ANY MASS STORAGE DEVICE.
```

It should be noted that we are liable to modify the use of the two types of mass storage from time to time, in order to run with the optimal operational environment, and any device dependent Mass Storage requests will eventually lead to applications failing.

- Tony Stanford

* * * * * * * * * *

NOS/BE 1.4 INTRODUCTION

During the past 4 or 5 months much systems programming effort has been spent in building and testing NOS/BE 1.4, which was then implemented in service on Wednesday 16 July 1980.

The primary motivation for performing this software upgrade is to provide operating systems support for the new 885 disc units and 7155 controllers which are to be installed in September 1980. The present NOS/BE 1.3 system does not support this hardware.

In addition, there have been many software bug fixes incorporated in NOS/BE 1.4 and also a few, relatively minor, software enhancements. Most of the software enhancements, however, are not relevant to our configuration or are relevant only to the machine operators, so that few user changes will be seen. However, one area of change is in the Advance Access Methods package of Record Manager, where additional facilities are available.

The software installed in service on July 16th is NOS/BE 1.4 PSR level 518. All the CDC products, with the exception of FTN and the Fortran Library were also updated to PSR level 518. FTN and the library will remain at level 473 (but see below). We expect to upgrade FTN and the library to level 518 towards the end of September 1980.

The following bug has been found in the present NOS/BE 1.3 system which is fixed in NOS/BE 1.4 and which may affect users. Record Manager under NOS/BE 1.3 allowed the use of a filename illegally specified as

```
6HFILNAM
```

such a filename will give RM ERROR 165 under NOS/BE 1.4 and should be correctly specified as

```
6LFILNAM
```

Under NOS/BE 1.4, the load map somewhat confusingly shows that Fortran library routines are from level 518. They are in fact from level 473, unless NEXT,FTN. has been used. (See below)

A bug in the level 473 Fortran Library allowed users illegally to issue a READ asking for more data than available in the existing record (i.e. READ list longer than record length). This bug is corrected in the level 518 Fortran Library and when this is implemented at the end of September, users at present taking advantage of this bug will be rewarded by a program abort of the form

```
LIST EXCEEDS DATA,FILENAME - lfn
FTN - FATAL ERROR 89
```

Programs should be altered to either READ the correct data length, or use BUFFER IN followed by the LENGTH function. Chapter 16 of the FORTRAN EXTENDED VERSION 4 REFERENCE MANUAL (60497800) describes the implementation of execution time input output.

If it is required to BUFFER IN a file lfn written with unformatted WRITE statements it is necessary to precede the program call by

```
FILE(lfn,RT=W,BT=I)
```

to redefine the file type.

The level 518 version of FTN, the Fortran Library and MANTRAP are available as test versions. To call this version of the compiler the job control statement

```
NEXT,FTN.
```

should follow the ACCOUNT card. NEXT,FTN will not work from interactive terminals. Please use this facility to check that your programs will not be affected by the new compiler when it is implemented at the end of September.

- Peter Gray

* *

THE MYSTERIES OF FORMATTED I/O

Formatted Read and Write statements handle coded data according to well-known Fortran conventions; so well known that one easily overlooks its innermost details, such as:

character code vs storage device;
data structure vs record management;
format specification vs CPU consumption.

These apparently harmless aspects of formatted I/O may have severe repercussions on two activities, viz:

a) data interchange between computers
and b) program execution.

A few rules about a) may easily be pointed out to readers who appreciate the difference between mass storage (PRU device) and SI,S/L tapes. These rules are based upon character code and data structure considerations. A hint for performance improvement can also be given about b) in terms of format specification.

a) Data interchange between computers

a.1) Cray coded data interchange

If data are to be sent to another Cyber/Cray computer complex a
DISPOSE,....,DF=TR.
to an SI tape will do. Otherwise you can only proceed by first converting to the Cyber coded data format. This can be obtained by using:
DISPOSE,....,DF=CB.
and remembering what follows.

a.2) Cyber coded data on disc

Data may have been written on a Cyber disc because of
formatted write;
DISPOSE,...,DF=CB. (from Cray);
job card deck submittal.

They are always, and only, represented in DISPLAY CODE, a six bit per character notation documented in Appendix A of the NOS/BE Reference Manual. Their structure can be described in Record Manager terminology as an RT=Z,BT=C type of structure (in other words, by the definition of their Record Type and Block Type). Detailed information about this can be found in the BAM (Basic Access Methods) User's Guide. Display code data in RT=Z,BT=C format can be transferred to any Cyber installation only. If this meets your requirements, then do a COPYCF of your disc file to an SI tape and off you go.

a.3) Cyber coded data on SI tape

Data may have been written to an SI tape by
formatted write;
DISPOSE,....,DF=CB. (from Cray);
COPYCF of any of the above (as from a.2).

These data are also in display code, RT=Z,BT=C format.
It should be noted that the tape blocks, in this case, are 128 Cyber words long (or 1280 six bit characters).
These tapes can be used at any other Cyber installation only.

a.4) Cyber coded data on SI "binary" tape

Data may have been written to an SI "binary" tape by
COPYBF (mistakenly or not) instead of COPYCF as from a.3;
formatted write, provided your program execution statement was preceded by
FILE,....,CM=NO.
where CM stands for Conversion Mode.

These data are still in display code, RT=Z,BT=C format but the tape blocks are now 512 Cyber words long (or 5120 six bit characters).
Tapes like this can also be used at any Cyber installation, subject to remembering how they were created, and therefore doing a COPYBF back to disc before processing their data, or formatted read, under control of a FILE,....,CM=NO. statement as above.

a.5) Cyber coded data on S/L tapes.

Data may have been written to an S/L tape by
formatted write, under control of a
FILE,...,RT=F,BT=K,FL=xx,RB=yy.
statement,
or
LISTCI/COPYCI (as from the ECLIB library documentation, soon to appear as Bulletin B6.1/3).

These data are either in EBCDIC or in ASCII code, depending on the presence or the absence of the parameter "EB" on the REQUEST control statement used to obtain the assignment of an S/L tape. These codes use 8-bits per character and the choice of one of them should be dictated by the "preference" of the computer receiving the tape. The structure of the data depends on the parameters specified in the FILE card, so that each K block (BT=K) contains as many fixed length records (RT=F) as indicated by RB=yy (Records per Block). Each record has a length of "FL=xx" characters. If some records need less than "xx" characters, trailing blanks to "xx" characters are used.

Note - No attempt should ever be made to use COPYCF from a disc file to an S/L tape. Those who might have done it should approach the Advisory Office to be enlightened.

b) Program execution

It is worth keeping in mind that the Fortran compiler stores the FORMAT almost unchanged, so that the entire interpretation has to be done at execution time, that is, for each Read/Write execution - so you can understand why so much more CP time is required by formatted I/O in comparison to unformatted. Still, one must be aware that the same line of coded information can be processed with considerable differences in CP time consumption, depending on the format ... of the FORMAT.

A significant example has been provided by some measurements performed at the ECN-NSP-REKENCENTRUM in the Netherlands, on a Cyber 175. The following has been taken from a presentation by the above installation made at one of the last conferences of CDC users.

A FORMAT (F7.3,10F10.6) was used to print out a table with 11 real numbers per line. Then, to be sure some spaces would always appear between the columns of the table, FORMAT(1X,F6.3,10(2X,F8.6)) was used. This simple change increased the CP time consumption by 65% !!!
A programmer can therefore perform his own optimisation of formats by adopting the following rules:

Use format specifications prefixed by a repetition count (e.g. 10F8.3) to transmit several variables by the same specification. In that case only one interpretation is needed to transmit several variables.
Include desired spaces in the field width of the conversion specification.
Avoid paranthesis if possible.
Remove redundant blanks and separators from variable execution time formats. This is not necessary for regular FORMAT statements in the program text, because then the compiler removes them.
Consider using an auxiliary array in the I/O list, because implied do loops of the form (R(I),I=1,IE) are transmitted faster than separate variables, compound implied do lists etc.
Use unformatted I/O for intermediate files.

- Luigi Bertuzzi

* * * * * * * * * *

VECTORISATION OF AN IRREGULAR INTERPOLATION PROBLEM

Post-processing the output from the ECMWF forecast model provides an interesting example of a problem which can be effectively treated by vector processing, although at first glance this appears difficult.

The basic task is one dimensional interpolation. A physical quantity, temperature for example, is known at several discrete values of a co-ordinate, (the sigma levels of the model), and an output temperature is desired at some intermediate position. If one was computing a single point only there would not be much scope for vector methods, but in fact a complete horizontal field is desired, and one can attempt to compute all the field values in parallel. With a normal interpolation problem this is straight-forward, and the horizontal loops can be moved inside for vector performance. However, the output temperature is specified in a different co-ordinate system, (pressure co-ordinates). Viewed in one co-ordinate system, the surfaces of constant elevation of the other system are not flat, but quite complex in shape. Although the co-ordinate transformation is vectorisable, the desired output points are pseudo-randomly placed relative to the data levels, and at different points in the horizontal, one needs data from different levels of the model. Consequently, one of the prerequisites for vector processing, (namely the need to access elements whose addresses form arithmetic progressions), is not fulfilled.

A "gathering" approach was used. A vector of integer pointer values was prepared. These values locate the output points relative to the model levels. (Since the values of the model levels may all be held in one vector register, it was possible to compute this index vector efficiently via function INDEX in ECLIB). Then, in several gather loops (of the form A(I)=B(IPOINT(I))), the irregularly located data needed were collected into consecutive locations. These loops are also implemented efficiently on the CRAY, via function GATHR in ECLIB. Now many points can be interpolated at once in a vector loop.

In the actual problem, further complications arose because of horizontally dependent extrapolation near the ground, with the extrapolation method depending on the physical quantity being interpolated. Consequently, this part was left in scalar mode. For temperature, the modified code took less than 25% of the CPU time of the original.

This approach requires several temporary vectors, and it is not, in general, possible to do complete fields all at once. In the problem under discussion, the code was already organised to do work one latitude circle at a time, so the storage penalty was not great.

The interpolation scheme was a cubic spline, but the method of vectorisation will work for other schemes. More complex schemes use more data, and hence do more gathering, but they do more computation, so that the relative gathering overhead does not increase.

The approach is easily generalised to higher dimensions, so that e.g. one could vectorise the interpolation to randomly located points on a horizontal grid of data (giving forecast values at weather stations is one possible usage).

- Dale Robertson

* *

ECMWF SUBROUTINE LIBRARY

The library ECLIB, which is available on both CYBER and CRAY, now contains the routines as listed in the last Newsletter (page 13, no. 3).

Documentation for these routines is being distributed as part of the Computer Bulletin set.

There is one addition to the previously published list of contents. The subroutine CONVAR interpolates between LAT/LONG and polar stereographic grids. It has been implemented in ECLIB with a small source change to overcome a problem associated with the SQRT and ACOS functions on the CRAY.

- David Dent

* * * * * * * * * * * *

* CYBER NAG LIBRARY - REMINDER

I have noticed that some jobs are still ATTACHing the permanent file NAGLIB. As the Mk.6 Library exists as a System Library, this is unnecessary. Also, the permanent file version is very old (Mk.5). Notice of this changeover was given in the October 1979 Newsletter. We intend to remove this version of the library on Wednesday 1st. October.

Also,the Mk.7 version of the library is available from permanent file, if anyone wishes to make use of it. The control cards needed for the two versions are as follows :

Current version	Mk.7 version
:	:
LDSET(LIB=..../NAGLIB)	ATTACH(NAGLIB,NAGMK7) LDSET(LIB=..../NAGLIB)

- John Greenaway

* * * * * * * * * * * *

USER SUPPORT STAFF CHANGES

S. Tinti has left the Centre, his position in User Support being taken by Michel Miqueu from Met. Applications section. One new member has joined the Centre, Norbert Kreitz.

As a new member Norbert Kreitz joined User Support on 1 July. He studied physics at Cologne and Düsseldorf University. He got his computer experience at the Düsseldorf University Computer Centre and especially at the Mülheim Max Planck Institut with PDP10. Before joining the Centre, he worked at an environmental protection institute run by the government of North-Rhine-Westphalia. As a system analyst he was involved in various problems using a CDC 6400 under NOS.

- Andrew Lea

* * * * * * * * * * * *

CYBER INPUT QUEUE DELAY STATISTICS

Some details of the input queue delay of jobs, depending upon their T and CM parameters, were produced for the TAC earlier this year. They are reproduced on the next page.

Though full up-to-date figures are not available, it is believed that some improvement has occurred in large CM jobs since the Cyber memory was increased.

- Gary Harding

* * * * * * * * * * *

Requested Time Limit secs.	Requested Maximum Memory Used x 1024 words	Average Input Queue Delay HH:MM	Maximum Input Queue Delay HH:MM	% No. of Jobs	% No. of Units
16	24	0:2	1:50	23.6	4.1
16	48	0:2	1:41	5.4	.6
16	104	0:20	5:35	3.2	1.4
64	24	0:21	48:18	35.6	10.5
64	48	0:31	8:58	6.7	4.4
64	104	3:3	144:27	11.5	6.6
256	24	0:26	2:41	1.9	2.2
256	48	2:36	48:18	2.0	4.7
256	104	7:54	65:43	4.9	15.7
512	24	0:51	5:5	.1	.2
512	48	0:8	1:40	.4	2.5
512	104	19:49	141:15	2.0	13.5
1536	24	0:19	1:2	.5	5.9
1536	48	1:35	11:16	.3	5.7
1536	104	17:33	86:1	1.8	21.9
1536	104	19:7	21:42	.1	.1

The total number of jobs run was 4037 and the total number of units charged was 46975.

As would be expected, jobs requiring more computer resources have to wait a longer time before starting execution. Jobs requiring less than 256 CP seconds and less than $24K_{10}$ memory are able to start execution quite quickly. However, once above this figure, delays become longer. Indeed, over 51% of the Cyber workload (by charge units consumed) has to wait an <u>average</u> of over 7½ hours before it begins execution, and over 35% of the workload is delayed by an average of over 17½ hours.

GRAPHICS

Old Graphics Libraries

The obsolete libraries VARIANLIB,ID=EWPLOT and NEWCONTLIB,ID=EWPLOT have now been removed from the system.

Old Coastline Database

The old coastline database BACKBASE,ID=EWPLOT will eventually be removed from the system. Although it is less accurate than the preferred database COASTLINES, it is still used on an occasional basis. For this reason, a copy has been made under the ID=PUBLIC. To use these databases with the contour package, use:

```
        ATTACH,TAPE90,BACKBASE.
or      ATTACH,TAPE90,COASTLINES.
```

No copy of BACKBASE exists on the Cray.

- Howard Watkins

* * * * * * * * * * * * * * * *

NOS/BE PERFORMANCE MONITORING PACKAGE

ECMWF have developed a software measurement package which can be used to evaluate the performance of the NOS/BE operating system. When using this package, the experienced NOS/BE systems programmer can identify bottlenecks in the configuration of the hardware and/or the operating system, and can then take steps to improve the overall performance of the NOS/BE system.

The package consists of a number of modifications to NOS/BE, plus a number of programs to analyse and reduce the data. Specifically, the programs consist of the following:

i) A method of identifying the most frequently called PPU programs. This will then allow the site to establish the best possible residency (Central Memory or Disc) for each PPU program.

ii) A method of analysing the transfer lengths of stack processor requests to the discs. This allowed ECMWF to analyse the relative performance expectations of 885 and 844 disc drives.

iii) A comprehensive statistics collection package, and graphical visualisation package which analyses:

- number of batch jobs completed per hour
- number of Intercom commands per minute
- % CPU utilisation
- CIO calls per second
- analysis of swapping activity for batch and interactive jobs
- memory utilisation by different job class.

The package at present works at level 473 and is being updated to run at level 518.

- Peter Gray

* * * * * * * * * * * * *

COMPUTER TRAINING COURSES

The Centre is offering another series of training courses for Member States personnel and ECMWF staff. Information has been sent to the Member States and nominations to attend are now invited. Nominations from ECMWF staff will shortly be invited via Section Heads.

In summary, the courses being offered are:

Course B BASIC USAGE 6 - 10 October

This is intended for anyone who will be actually programming the Cray-1, to give them sufficient experience to run simple work. It will also introduce them to some of the Cyber facilities they may need to complement their Cray activity. Prior knowledge of another computing system, plus a knowledge of Fortran, is required. An optional 5th day (10 October) is devoted to explaining how to use ECMWF's meteorological database and archive system.

Course C CRAY USER COURSE 13 - 17 October

An in-depth course for those who will make heavy use of the Cray-1 and its many unique facilities. Intending participants will be expected to know how to run simple jobs on the Cray-1.

Course D CYBER USER COURSE 27 - 31 October

Again, an in-depth course, this time for those who will make detailed use of the Cyber 175 in conjunction with the Cray-1. Among other things, it will cover Cyber file handling (including private tapes and discs), and program preparation facilities (UPDATE, debugging, etc.). Intending participants will be expected to know the basics of a Cyber NOS/BE system.

It is planned to offer a similar set of courses in 1981.

- Andrew Lea

* * * * * * * * * * * * * *

DOCUMENTATION

CDC manuals to bring documentation up to PSR level 518, which is the recently installed version of the operating system, are being distributed. Various bulletins, including ECLIB documentation, will also be distributed shortly, and will be accompanied by a new 'Contents' sheet, so that users will be able to ensure that they have received a copy of all available bulletins, and are in possession of the most recent version.

- Pam Prior

* * * * * * * * * * * * * *

MEETING OF MEMBER STATE COMPUTING REPRESENTATIVES

A meeting of Member State Computing Representatives will be held at the Centre from 21 - 23 October. Invitations to Member States to nominate attendees have been sent out.

The idea of such a meeting is to enhance the informal contacts between Centre staff and Member State users of the Centre's computing facilities. This should help Member States to use the Centre's facilities more effectively, and allow the Centre more fully to understand Member State day to day problems, thus assisting the Centre in providing a more effective service.

- Andrew Lea

* * * * * * * * * * * * * *

STILL VALID NEWS SHEETS

Below is a list of News Sheets that still contain some valid information which has not been incorporated into the Bulletin set (up to News Sheet 87). All other News Sheets are redundant and can be thrown away.

No.	Still valid article
11	FTN Rounding Option
15	Private Packs on the Cyber (MOUNT/DISMOUNT)
16	Checkpointing and program termination
17	Private packs and interactive jobs
19	CRAY UPDATE (temporary datasets used)
31	Fortran Callable Tape REQUEST
37	IN trays for Cray and Cyber jobs
42	Cyber Scheduler (see News Sheet 59 also)
43	Cray AUDIT
	Transfer of Coded Files
47	Libraries on the Cray-1
50	8 Disc Cray System
	Terminal Procedure
51	Cyber Disc Reconfiguration
53	Cyber Job Card Priority Usage
	Writing 6250 bpi Tapes (EEC Parameter)
	Punching Conventions (Coding Forms)
54	Things not to do to the Station
55	New Cyber Peripherals
56	DISP
59	New Cyber System (Scheduler Changes)
64	New Version of Graphics Software
65	Data Security on Cyber and Cray
66	New Cray Audit
	Cyber Accounting
67	Attention Cyber BUFFER IN Users
70	Cyber/Cray Station
71	Packs Command
72	The Change to BST (Machine Schedules)
73	Minimum Cyber Field Length
75	Disposing with SDN=PLOT
76	New Version of Graphics Software
77	ACCOUNT of an Executing Job
80	Fortran READs at Levels 473 and 508
86	NOS/BE 1.4 Introduction
87	New Graphics Software - Version 110

The News Sheets which can be thrown away since this list was last published are numbers 63, 68, 78, 79, 81-85.

- Andrew Lea

* * * * * * * * * * *

MEMBER STATES' USAGE OF CRAY RESOURCES UP TO JULY 21 1980 (IN UNITS)

Denmark	5965 units
Finland	163 units
France	39495 units
Germany	1 unit
United Kingdom	95556 units
Netherlands	1997 units
Spain	3341 units
Sweden	14802 units
Yugoslavia	1916 units

- Andrew Lea

* * * * * * * * * * *

INDEX OF STILL VALID NEWSLETTER ARTICLES

This is an index of the major articles published in the ECMWF Newsletter plus those in the original ECMWF Technical Newsletter series. As one goes back in time, some points in these articles may have been superseded. When in doubt, contact the author, or User Support.

GENERAL

METEOROLOGY

* T indicates the original Technical Newsletter series.

USEFUL NAMES AND 'PHONE NUMBERS WITHIN ECMWF

		Room*	Ext.**
Head of Operations Department	- Daniel Söderman	OB 010A	373
ADVISORY OFFICE - Open 9-12, 14-17 daily Other methods of quick contact: - telex (no. 847908) - COMFILE (see Bulletin B1.5/1)		CB 037	308/309
Computer Division Head	- Geerd Hoffmann	OB 009A	340/342
COMPUTER OPERATIONS			
Console	- Shift Leaders	CB Hall	334
Reception Counter) Terminal Queries)	- Judy Herring	CB Hall	332
Operations Section Head	- Eric Walton	OB 002	349/351
Deputy Ops. Section Head	- Graham Holt	CB 033	334
DOCUMENTATION	- Pam Prior	OB 016	355
Libraries (ECMWF, NAG, CERN, etc.)	- John Greenaway	OB 017	354
METEOROLOGICAL DIVISION			
Division Head	- Roger Newson	OB 008	343
Applications Section Head	- Joel Martellet	OB 011	360
Operations Section Head	- Austin Woods	OB 107	406
Meteorological Analysts	- Ove Åkesson	OB 106	380
	- Veli Akyildiz	OB 104A	379
	- Horst Böttger	OB 104A	378
	- Pauno Nieminen	OB 104A	378
	- Herbert Pümpel	OB 106	380
Meteorological Operations Room		CB Hall	328/443
REGISTRATION (User and Project Identifiers, INTERCOM)	- Pam Prior	OB 016	355
Research Department Computer Co-ordinator	- Rex Gibson	OB 126	384
Systems Software Section Head	- Peter Gray	CB 133	323
Tape Requests	- Pauline Litchfield - George Stone	CB Hall	332
TELECOMMUNICATIONS			
Fault Reporting	- Pierre-Pascal Regnault	CB 028	397/375
Section Head	- Fritz Königshofer	CB 130	310
User Support Section Head	- Andrew Lea	OB 003	348

---

* CB - Computer Block
OB - Office Block

** The ECMWF telephone number is READING (0734) 85411, international +44 734 85411